

AFRL-RI-RS-TR-2009-151
In-House Final Technical Report
June 2009



ASSURED COMMUNICATIONS RESEARCH CENTER (ACRC)

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2009-151 HAS BEEN REVIEWED AND IS APPROVED FOR
PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION
STATEMENT.

FOR THE DIRECTOR:

/s/

DAVID L. BIBIGHAUS, Maj., USAF
Chief, Cyber Defense Branch

/s/

WARREN H. DEBANY Jr.,
Technical Advisor Information Grid Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE*Form Approved*
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) JUN 09		2. REPORT TYPE Final		3. DATES COVERED (From - To) Dec 00 – Dec 08	
4. TITLE AND SUBTITLE ASSURED COMMUNICATIONS RESEARCH CENTER (ACRC)				5a. CONTRACT NUMBER In House	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER 62702F	
6. AUTHOR(S) Kevin A. Kwiat and Kathaleya Wolfe				5d. PROJECT NUMBER 4519	
				5e. TASK NUMBER 22	
				5f. WORK UNIT NUMBER 49	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFRL/RIGA 525 Brooks Rd. Rome NY 13441-4505				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/RIGA 525 Brooks Rd. Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) N/A	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2009-151	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 88ABW-2009-0904 Date Cleared: Mar. 2009					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In 2000, Air Force Research Laboratory, Information Directorate, Rome Research Site, established the Assured Communications Research Center (ACRC). The program was designed to promote and perform extensive research in the transformation of concepts that originated in the domain of fault tolerance to the then emerging area of information assurance. When a computing technique is placed in an information assurance scenario, the technique must be guarded from an attacker. A basic tenet of the ACRC was to enforce the system properties of fault-tolerant computing even in the presence of attack in order to ensure system integrity and availability.					
15. SUBJECT TERMS Assured Communications, Fault Tolerance, Voting, Intrusion Detection, Coordination Models, QoS, Loss Inference, Wireless.					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 39	19a. NAME OF RESPONSIBLE PERSON Kevin A. Kwiat
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

Table of Contents

1. Introduction.....	1
2. ACRC Foundational Work: Distributed Voting	3
3. Adapting Distributed Voting Algorithms	6
5. Collaborations: Analyzing Intrusion Detection Systems	10
6. User-Level Intrusion Detection.....	13
8. Simulation-Based Validations	16
9. Secure Knowledge Management	17
10. Coordination Models	18
11. Predicting Real-Time Data Limits	19
12. Updating SCADA Systems through External Coordination.....	21
13. Graphical User Interface (GUI) Based Systems	22
16. Overlaid Wireless Access Networks in Market-Based Environment	28
17. Maximizing Spectrum Utilization in IEEE 802.22	29
18. Education for Potential Air Force Officers	30
19. Summary	31
20. References	32

List of Figures

Figure 1: ACRC Research Timeline	2
Figure 2: Majority Voting Architecture	4
Figure 3: Concurrent Intrusion Detection System	11

1. Introduction

When an attacker in cyberspace has out-maneuvered our ability to prevent, avoid, and detect the attack, then the attack successfully breaches our cyber defenses. Air Force systems must survive and recover. Such cyber attacks may cause unforeseen damage to our systems, so we must have a dynamic capability to survive and recover from the damage. However, survival and recovery must emphasize maintenance of the overall mission capability rather than on the survival of individual hosts. This ability renders our systems with the ability to “fight through” an attack.

Survivability is the quantified ability of a system, subsystem, equipment, process, or procedure to continue to function during and after a natural or man-made disturbance. For the Air Force to establish information dominance, its systems must be survivable by possessing the ability to tolerate errors or attacks and gracefully degrade with respect to mission critical requirements. Systems may also be subjected to constant change such as overload, component failure, evolving operational requirements, and a dynamic operational environment. Therefore, the Air Force has a critical need for techniques and technologies that make information systems survivable. Survivability allows Air Force systems to absorb the punishment of an attacker and have the foundation of system reserves to fight through the attack.

Deployment-time issues almost always dictate that information system survivability be carried out automatically. Otherwise, the associated time-lag with manual intervention would render the system as unavailable for too long. Furthermore, the ability to “fight through” implies the need to “continue through” with the system’s intended – that is, correct - functionality. The automatic triggering of mechanisms to ensure availability and the concurrent maintaining of correct system functionality are closely related to two basic attributes of information assurance: availability and integrity. New security advancements are continuously being made; nevertheless, although these may focus on another information assurance attribute of confidentiality, authentication, or non-repudiation, each advancement should also consider availability and integrity. If availability and integrity issues are not addressed they can lead to a false sense of security because a situation could arise where

attackers target the security mechanisms themselves. When engaged in an attack, the security mechanisms can be either made ineffectual or deliberately evaded. The Assured Communications Research Center (ACRC) took these matters into consideration by adopting a broad view of communications: in striving for network information assurance, the ACRC also considered communications between components that are within a single computer – especially communication between redundant components that comprise a fault-tolerant system. In adopting a broad view of communications, the ACRC collaborated with researchers that represent the following academic institutions: University at Buffalo, University of Central Florida, Illinois Institute of Technology, Stevens Institute of Technology, Union College, Hamilton College and the City University of New York. Figure 1 shows the various research tasks undertaken by the ACRC.

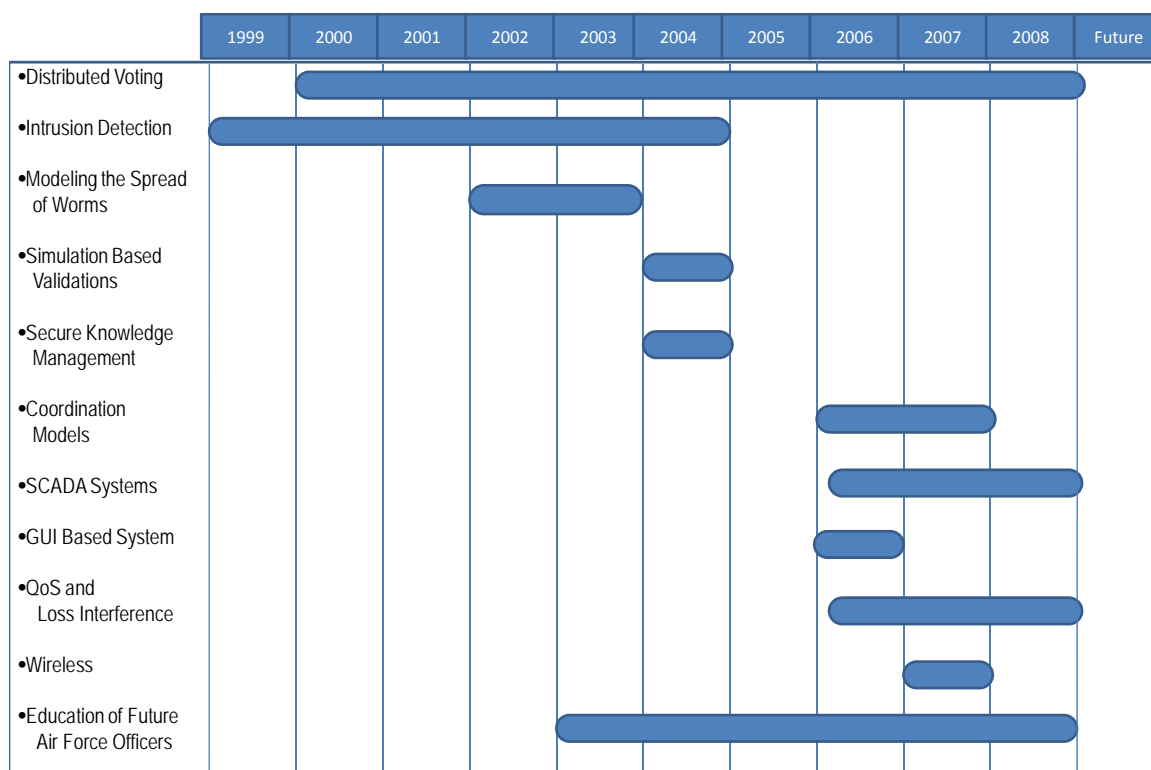


Figure 1: ACRC Research Timeline

2. ACRC Foundational Work: Distributed Voting

The combined attributes of survivability and continuity can be sustained with functional replication. These replicas can be considered redundant; yet a fault-tolerant system possesses some redundancy so that it can withstand internal failure. Conversely, a system lacking seeming “extra” components cannot be expected to persevere when the simplex component malfunctions. Adding redundancy alone, however, is never enough; instead, some system features are needed to coordinate the redundancy so that it makes the system fault-tolerant. One such feature is voting. Voting is performed on the outputs of redundant components to produce a single error-free output that represents the majority of the redundant components. Voting is therefore seen as a pervasive element for a “fighting through attack” capability. The Assured Computing Research Center began its research in 2000 with the technical analysis of distributed voting algorithms. Dr. Kevin Kwiat and Ben Hardekopf of the ACRC published an article entitled “Performance Analysis of an Enhanced Security Distributed Voting Algorithm” in the SPECTS 2000 Symposium. Their goal was to develop an algorithm that would surpass the current fault-tolerant system designs. Distributed voting algorithms utilized the predominate two-phase commit protocol, where vote comparisons were made on each processor in a system, then the sharing of votes took place to determine majority vote. Once majority vote was reached, it was committed. To construct the algorithm, Kwiat and Hardekopf took a non-intuitive approach of reversing the two-phase commit protocol. Their design committed first, and decided on majority vote. Their analysis and simulation results demonstrated how their algorithm suffers zero loss in fault coverage while simultaneously yielding both a security benefit and a gain in performance [1]. Figure 1 illustrates the voting architecture.

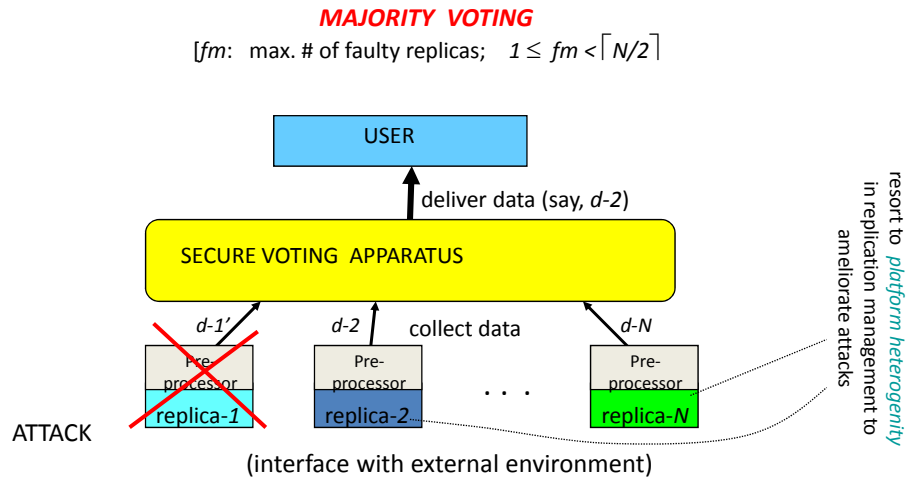


Figure 2: Majority Voting Architecture

One example of a distributed voting algorithm that utilizes the two-phase commit protocol is known as DVA1 in the paper Kwiatt and Hardekopf compiled. DVA1 is a representation of the state-of-the-art voting algorithms, and is utilized in comparison to the algorithm applied by Kwiatt and Hardekopf. DVA1 provides a simple approach to voting. Once a task has been given, each processor in DVA1 calculates its result and broadcasts it to all other processors. Each processor then analyzes the votes. Majority vote is then taken to be the final result, and a random processor is chosen to send this result to the interface module. Any processor vote not included in the majority is considered faulty, and must initiate a recovery standard. If a majority vote is not reached, each processor runs self-diagnostic tests. Once processors have been determined as faulty or not, a non-faulty processor is chosen to send its results to the interface module. [1]

The essential change in the architecture of DVA1 to DVA2 is that the interface module of DVA2 contains a timer, which begins as soon as a value is committed. Once a task is sent out, each processor calculates its result. One processor is then arbitrarily chosen to commit the result to the interface module, starting the timer. Once the other processors view the result, they compare it to their own votes, and broadcast a dissenting vote if they do not agree. If majority processors do not agree, a new vote is sent to the module for comparison. The module discards the old vote, and restarts the timer. This process is repeated until majority processors agree with the vote and it is sent on to the user. Simulations of DVA1 and DVA2 show that no loss in fault coverage is associated with DVA2, even with the number of entries set to only one vote. The security benefit of this algorithm comes from the fact that even if an attacker manages to compromise a system and reside between the user and the coordinator, an incorrect result cannot be committed without the uncompromised processors correcting it. An attacker would have to compromise the majority of the processors before attaining the ability to submit an incorrect result. DVA2 monitors the output of the module, when the coordinator commits the final result. Errors in that stage are caught and corrected that otherwise would have passed through to the user in DVA1. Although holding results in a timed buffer increases time delay, DVA2 also produces a better performance gain than that of DVA1. Therefore, DVA2 improves both security and fault-tolerance levels in distributed voting [1].

Together in 2001, the ACRC published a technical report [2] based on the DVA2 algorithm that had been more aptly renamed as TBDVA – Timed Buffer Distributed Voting Algorithm. Kwiat and Hardekopf focused this report on secure distributed voting algorithms to analyze protocols and indicate weaknesses in cyber security. There have been various systems proposed with the purpose of making distributed voting more secure. A common weakness among these algorithms includes the fact that these systems utilize exact voting, and do not consider the requirements for inexact voting. Kwiat and Hardekopf describe in detail a protocol combining the requirements of security, fault-tolerance, and performance while remaining general enough to handle both voting procedures [2]. Their research in distributed voting algorithms led to both Kwiat and Hardekopf earning a United States Patent.

3. Adapting Distributed Voting Algorithms

Voting is a simple and effective way of managing replicated data in distributed systems [3]. In 2004, the ACRC hired two summer interns to complete research in Distributed Voting Algorithms. These interns, Dr. Kaliappa Ravindran and Ali Sabbir of the University at Buffalo, spent 8 weeks working with Dr. Kevin Kwiat of the ACRC on a paper entitled “Adapting Distributed Voting Algorithms for Secure Real-Time Embedded Systems” [3]. Their goal was to develop machinery that dynamically adjusted its internal mechanisms to deal with various types of failures in infrastructural levels. Information assurance goals pose a variety of threats because failure behaviors such as data corruptions and message timeliness violations in networks occur often. These threats occur even more so in wireless network settings. Their goal was to reduce message overhead and decrease power drain in wireless networks to ensure the absence of these threats. Kwiat, Ravindran and Sabbir focused their research paper on voting protocol mechanisms that deal with data corruption and message timeliness violations. They wanted to employ a variant of the two-phase majority voting protocol to validate data by processes. These protocol mechanisms employed multicast systems in the transport network. Underlying transport networks often provide little or no guarantee of communication reliability. Their protocol was designed to deal with these failures in system management to the extent required to establish majority consent or dissent vote on data [3].

Kwiat, Ravindran and Sabbir developed their algorithm utilizing multicast semantics, voter silence to consent to votes, and channel attacks for intricate detection. Their developments satisfied the safety conditions for voting protocols. Corrupted data or data that had reached its time bound is never transmitted to the end-user, even under strenuous failure scenarios. The three researchers concluded that protocols strive to satisfy the liveness condition to deliver uncorrupted data to the end-user with the time allotted. Their voting protocol mechanisms adapted to failure conditions occurring at the infrastructural level to meet information assurance goals [3].

In a distributed embedded system, data items can be shifted from one set of functional modules to another. This occurs through a common buffer interposed between the modules over a computer network [4]. To allow a timely processing of data by application modules, Kwiat, Ravindran and Sabbir described a programming structure based on the time-atomic write method. Entitled “Incorporating Timeliness in Atomic Write of Shared Data for Distributed Embedded Systems,” the paper described the structure and semantics of time-atomic write based on a notion of epochs [4]. An epoch is a time-interval crucial to the application over which the computation modules reach agreement on processed data items. These modules extracted a consistent view of shared data items processed and met the deadlines associated with these items. The time-atomic write primitive joins the flow of real-time events with the notion of reaching an agreement over shared data in distributed systems. Their paper established a means for coordinating an application’s distributed processes in a real-time environment.

4. Continued Research in Voting Algorithms

Real-time data collection systems in sensor networks are often faced with failures occurring during the data collection. The exposure of sensor nodes in hostile conditions often exhibits as data corruptions by malicious devices and timeliness violations in the processing and communication paths [5].

At the 2008 International Conference on Availability, Security and Reliability, Jiang Wu, from the City University of New York, K. Ravindran, Sabbir and Kwiatt presented a paper entitled “Adaptive Voting Algorithms for Reliable Dissemination of Data in Sensor Networks” [5]. The focus of this paper is to deal with data corruption and enforce timely delivery of correct data to the end-user by adding extensions to the voting protocol mechanisms. They describe an extension to the voting protocol mechanism that dynamically adapts to deal with various types of failures.

The team describes two modified versions of the two-phase voting protocol: an explicit and implicit voting mode. The latter is to enhance performance under normal circumstances. The explicit two-phase voting protocol known as M2PC or modified two-phase commit voting, requires a YES or NO vote from each node to a central vote collator. Only votes received are considered in the decisions. This produced a large amount of overhead and thus raises scalability concerns.

The implicit-consent voting protocol requires only dissent votes. This technique takes longer to complete the voting because it waits the full time-out period to determine dissent votes. However, this method reduces overhead and solves scalability concerns. The TBDVA-derived protocol known as IC-M2PC (implicit consent - modified two-phase commit based voting), rests on the foundation that the majority of voters are non-faulty and message loss is sporadic.

The augmentations to the protocols increase their robustness and performance. The important goal is to ensure the integrity of data delivery to the end-use in the presence of data corruptions and other faults in the sensing system. This protocol is highly adaptive to the various types of failures that would otherwise from data integrity attacks emanating from compromised voters.

5. Collaborations: Analyzing Intrusion Detection Systems

In 2000, Dr. Shambhu Upadhyaya from the State University of New York at Buffalo joined the ACRC to conduct research with Dr. Kevin Kwiat on intrusion detection systems (IDSs) in distributed information systems. The two researchers published a paper entitled, “A Comprehensive Simulation Platform for Intrusion Detection in Distributed Systems,” [6] that describes the simulation of an attack recognition system in a distributed environment. The technique of attack recognition is based on assertion testing. In their analysis they describe how an auxiliary process known as Watchdog queries a user for scope-files, from which an assertion plan called Sprint provides a platform for attack monitoring. Testing and evaluation environments for the intrusion detection system are generated to simulate this platform. Different attack scenarios are also applied to test the abilities of the recognition system [6].

The team found that this simulation performed well within this particular university environment because it was observed that normal computer sessions had six to eight processes running on the a computer system at a given time. The intrusion detection system they designed was able to detect all intrusive activities and terminate the connection for all log-ins of intrusive users in their test simulations. They concluded online intrusion detection that utilizes assertion testing provides low performance overhead and high-quality detection coverage. The system generated few false positives that occurred when a normal user selected minimal false entities from a session scope [6]. Figure 1 displays the basic block diagram of the team’s concurrent intrusion detection system.

Basic Principle




Figure 3: Concurrent Intrusion Detection System

In 1999 Kwiat and Upadhyaya examined whether monitoring user commands verses network traffic would produce more effective monitoring. Along with computer engineering student Ramkumar Chinchani from the University at Buffalo and the support of ACRC, they developed a software program in 2002 which generated profiles to guard against criminals at borderlines and ports. The user-level anomaly detection software demonstrated higher levels of protection for military and government installations. It also provided protection for banking and other commercial networks targeted in such attacks. The program designed updated profiles by tracking the performance of a user as they performed an array of routine tasks. By closely monitoring these actions, the program depicted malicious activity through a controller's application of email, archive searching, and file opening [7].

Their software program was designed to surpass other computer-security products that featured user-profiling techniques. They concluded the competitor software was generally 60-80% reliable, whereas the software they created would be up 94% reliable in these detections. Although proven effective, this software program would be one of thousands of tools in a computer-security arena requiring multi-layered defenses, Upadhyaya said [7]. Intrusion detection often results in false alarms, and keeping these alarms to a minimum was a key driver in this software's development. The overall success of this security software was featured in the magazine, *Scientific American*, in 2002 [8].

2003 launched a new year for intrusion detection research. On 24 March 2003, the annual IEEE International Information Assurance Workshop took place in Darmstadt, Germany. Here Chinchani, Kwiatt, and Upadhyaya introduced a paper on detection of attacks in user space. The paper, entitled "A Tamper-Resistant Framework for Unambiguous Detection of Attacks in User Space Using Process Monitors," [9] provided an alternate mechanism for unambiguously monitoring user space daemons. The paradigm generated profiles for every user in a system where security was a concern. They speculated such a framework must itself be tamper-resistant to attacks. The framework was based on concepts in topological fault-tolerance, and showed it could unambiguously detect such attacks yet provided a defense that would be very difficult to destabilize.

Given process monitors in known fault-tolerant arrangements, the team realized it was possible to disable services without being detected. Techniques in fault-tolerance applied prior to this held high inherent differences between faults and attacks. This proved they were not as sufficient as the theoretical framework the team introduced. The drawback to such an advanced structure was that operating systems at the time provided minimal support for the implementation in terms of asynchronous event monitoring. The framework could only be deployed when the installation system had not been compromised [9]. The ACRC did not view this as a failure, but as a successful attempt to shut down services on a process monitor without detection. This framework would prove very effective in fault-tolerant arrangements. Not only did it provide a method to bypass detection, it also generated techniques to shut down an attacker's architecture without their acknowledgement.

6. User-Level Intrusion Detection

In 2005, the focus of the research by Kwiat and Upadhyaya included University at Buffalo (UB) doctoral student Ramkumar Chinchani. Chinchani completed research with the Kwiat and Upadhyaya in 2002 and 2003 on Tamper-Resistant Frameworks. Chinchani published a doctoral dissertation entitled *A Job-Centric Approach to User-Level Intrusion Detection*, in which he addressed multiple issues concerning user-level intrusion detection and user-level threats [10]. He demonstrated his proposed solutions in the form of figures, graphs, architectures, and codes. He framed his research results so that mainstream computer networking might readily adopt his ideas on intrusion detection systems. With the support of the ACRC and the Computer Engineering Department at UB, Chinchani successfully completed his dissertation later that year.

Chinchani's thesis accentuated the difficulties encountered when designing an intrusion detection system based on anomalies as opposed to misuse. Anomaly detection relies upon finding a sufficient number of commands to warrant declaring a user's behavior as malicious; whereas misuse detection is when a user's behavior maps into a pattern that is known to be malicious. To make anomaly detection more succinct and tractable, a new approach was taken. . A higher order representation of a user profile was proposed, where the system documents the steps taken by the user to properly carry out commands. This would ensure user involvement in the security process and lower the rates of false positives. Chinchani also suggested *key challenge graphs*, which target user-level threat modeling. Threat models are utilized tools designed to provide a viable alternative in testing on live networks. If used appropriately, this technique would eventually eliminate threat models such as attack graphs and privilege graphs [10].

In 2005, the intrusion detection work reached another milestone when Mr. Chinchani defended his doctoral dissertation with Dr. Upadhyaya serving as his advisor and Dr. Kwiat as a member of the defense committee. That same year, the three became co-authors for a book chapter [11]. Dr. Chinchani's dissertation, in addition to documenting the issues concerning user-level intrusion detection, included several advancements such as a proposed higher-order representation of a user profile where the system documents the steps taken by the user to properly carry out commands. This would ensure user involvement in the security process to further lower the rates of false positives.

7. Modeling the Spread of Worms

The foundational work of the ACRC in replica voting as a basis for a “fight through” capability had led to question the validity of assuming that a sufficiently large number of replicas could withstand a fast-spreading attack. Whereas the intrusion detection work lent insight on how a replica might defend itself from attacks emanating from outside the voting network, attacks from compromised voters upon other voters needed to be characterized. Internet worms embody fast-spreading attacks that exploit common vulnerabilities of the attack’s targets. Dr. Lixin Gao of the University of Massachusetts, Amherst, joined the ACRC in 2002-2003 to analyze the malicious scanning methods employed by worms. The analysis revealed the resources worms need to perform the scans that then enable the worms to spread. The most recent worms at that time used, with dramatic effect and efficiency, some combination of scanning methods. From the ACRC’s efforts came a generic worm detection architecture that monitors malicious activities by detecting the spread of worms using real-time traces and simulations. The published results [12, 13, 14] shed light on how to defend against worms.

8. Simulation-Based Validations

The ACRC continued its work with Dr. Kaliappa Ravindran in 2004 by providing a simulation-based methodology to validate the correctness of a candidate protocol. Protocol correctness in concurrent systems describes the achievement of a protocol to its intended goal without any indication of faultiness on an application [15]. Along with City University at New York student G. Ding, their methodology was based on protocol state machine modeling in context with its operating environment. They published a paper [13] that offered an approach to a validation model that captures the external constraints at the protocol interface, the motion disturbances, and the internal rules and procedures for the protocol. Discrete simulators developed in suitable modeling language (such as PROMELA) verify safety and liveness conditions of a protocol. These conditions are prescribed at the external interface level under a given operating environment. This methodology demonstrates validation exercises on a two-phase commit based majority voting protocol [15].

The correctness of a protocol offers a set goal guaranteed to be achieved in all aspects of protocol execution. These achievements are executed without causing faulty effects on the application. The issue of correctness becomes more vital in Information Assurance employments, given a protocol is designed to deal with a hostile and unpredictable process behaviors such as Denial of Service (DoS) attacks. Kwiat, Ravindran and Ding offer ideas in their validation approach that model the protocol environment and interface. These ideas include the application of distinct Codesign Finite State Machines (CFSM), and then analyzing them in conjunction with a unique CFSM that represents the protocol itself [15].

9. Secure Knowledge Management

In 2004, the AFRL Information Directorate Contracting Division granted sponsorship to the 1st Workshop on Secure Knowledge Management that was held at the University of Buffalo (UB) in New York. The subsequent workshop support to UB and organizational involvement of the ACRC were deemed essential toward raising the awareness of academics and practitioners in secure knowledge management - a critical area of research. UB was selected by the National Security Agency as a Center of Academic Excellence in Information Assurance Education. Directing this program was Dr. Shambhu Upadhyaya. Dr. Kwiat explained how AFRL/IF and the ACRC played a major role in Dr. Upadhyaya's contributions to Secure Knowledge Management. Dr. Upadhyaya gained recognition at the University of Buffalo by highlighting several years experience at AFRL/IF as a National Research Council Summer Faculty Fellow and as an AFRL Visiting Faculty Member.

The researchers ensured Knowledge Management Systems (KMS) promoted the sharing of information to increase productivity; however the US government and other organizations had amplified concerns involving KMS. Kwiat stressed the growing concern of web-access and intranets, stating the high importance of the protection necessary to secure corporate knowledge as companies continue to give access to numerous individuals. Dr. Kwiat emphasized how AFRL and the ACRC aided in the research for securing information, data, and intelligence through Knowledge Management Systems. Due to previous work with the ACRC and UB, the workshop became a baseline in guiding students to join the NSA-accredited Information Assurance Center to undertake challenging research endeavors that were critical to the work at AFRL/IF. The workshop provided gateways for further investigations and projects that were important to the ACRC.

10. Coordination Models

In 2006, Dr. Kwiat performed research with Dr. Shangping Ren from the Illinois Institute of Technology who joined the ACRC as part of the AFOSR-sponsored American Society of Engineering Education-administered Summer Faculty Fellowship Program. The two researchers presented a paper at the 2006 IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC). The paper, entitled “A Coordination Model for Improving Software System Attack-Tolerance and Survivability in Open Hostile Environments” [16], describes a proposed coordination model that contained three active entities: actors, roles, and coordinators. These three entities compose what they referred to as the ARC model. The actors abstracted the system functionalities, while the roles and coordinators statically summarized the system’s constraints. These constraints were then broadcasted and passed among themselves and onto the actors. Kwiat and Ren noted that a software system’s attack-tolerance and survivability in open hostile environments are enhanced through suitable propagations constraints and their enforcements [16].

The two individuals discussed the properties of their proposed model. They presented their ARC model in a way where the security and fault-tolerance policies were transmitted among mutual check components. They also described multiple survivable feedback loops built in the ARC model and discussed how the model self-healed from attacks on individual portions of the ARC. Kwiat and Ren attempted to utilize coordination strategies to improve system survivability. Their model not only avoided a single point of failure in coordination for this environment, it also provided two survivable feedback loops that detected possible attacks. These loops transformed fault-tolerance techniques to attack-tolerance tools in a coordinated model approach. Along with Yue Yu from the Illinois Institute of Technology and Jeffrey Tsai from the University of Illinois at Chicago, Drs. Kwiat and Ren published an extended version of their SUTC paper in a 2007 issue of the *International Journal of Distributed Sensor Networks* [17].

11. Predicting Real-Time Data Limits

In a real-time environment, the semantics and the importance of data depend on the time the data is used. The process of getting a consensus data from a group of replicated units must not take longer than the life expectancy of the data. In a paper entitled “Take Intelligent Risk and Optimize Decision Based on Time, Available Resources and Risk Tolerance Limits” [18], Dr. Kwiat and Dr. Ren presented analytical solutions for the expected time when dependable data could be obtained under unique voting schemes [18]. They also demonstrated how the allocation of resource played a significant role in satisfying both data accuracy and consistency constraints.

For real-time data, the semantics and importance of data are largely time-dependent. Therefore, real-time data typically has a lifespan associated with it. Most voting schemes in real-time settings must apply deadlines to mark the end of a data lifespan. If the deadline is reached before the resultant consensus is acquired, the data is sacrificed and a new round of data solicitation is commenced. Drs. Kwiat and Ren discussed the expected time to obtain a valid vote and explained the concept of interval-based timing constraint models. They also demonstrated that when point-based constraints are intrinsically impossible to satisfy, a more general interval-based constraint could be used to obtain satisfactory solutions. The two researchers concluded that, with statistical assumptions, an increase in replica voter numbers reduced the expected time to obtain assured or trustworthy votes. Although doing so increased the trustworthiness of the voting system, it offered little improvement on the delivery time expected for the voting process.

The ACRC researchers also focused research on vulnerabilities that exist in Supervisory Control and Data Acquisition (SCADA) systems. Together with Ren’s graduate students Kun Xiao, Nianen Chen, Limin Shen, and Xianhe Sun from the Illinois Institute of Technology and Michael Macalik from the Rome Research Corporation, they published a paper detailing a non-intrusive approach for improving the survivability of SCADA systems without interruption of their normal process flow [19]. SCADA systems avoid unsafe conditions by including interlocking logic code on the base system. This prevents conflicting operations from starting at inappropriate times, and provides corrective action or shutdown of the system when an unsafe condition is detected. The team proposed a workflow solution that was constructed on a system outside of

the attack path and separate from the process under control. The solution consisted of a simulation of the SCADA system's workflow. The cause-and-effect relationship of a set of commands that were to be processed by the SCADA system were first performed, in simulation, to help detect malicious operations before the effects of such operations would manifest on the actual SCADA system. . For this work, the researchers created a notional model of a SCADA system and used only open-source information in formulating the model. The model's workflow stored functional and survivability knowledge about the system and simulated a fault whenever failures induced by malicious logic were executed. The modeling of these modes of failure proved valuable in implementing damage control. This model was event driven and conducting the simulation externally prevented interference with the normal functionality of the underlying SCADA system. The "separation of concern" principle was reflected in this model's treatment of security and survivability concerns apart from those concerns that deal with supervisory control and data acquisition. This may enable the team to accommodate future requirements that could not have been anticipated due to continually evolving threats.

12. Updating SCADA Systems through External Coordination

Kun Xiao, Shangping Ren and Kevin Kwiat's continued research on SCADA systems revealed significant problems with updating these types of systems. Their 24x7 availability requirement prevents the traditional "shutdown and update" approach. The three individuals published and presented a paper on this topic entitled "Retrofitting Cyber Physical Systems for Survivability through External Coordination" in 2008 at the 41st Hawaii International Conference on System Sciences [20]. The paper proposed a new approach to upgrading legacy SCADA systems. This method constructs an external coordination layer, separated from the process under control, which only interfaces with the SCADA systems through events. The coordination layer combines the fault-tolerant schemes of the SCADA systems and a model to coordinate critical services when cyber attack occurs. This approach greatly reduces the disturbance to the underlying system because the survivability-related knowledge and protection scheme are built in the coordination layer that is external to the SCADA systems. The second advantage is the separation of fault-tolerance and survivability concerns from supervisory and acquisition which enables accommodation for future requirements [20].

13. Graphical User Interface (GUI) Based Systems

In mid 2006 Kwiat, Upadhyaya, and two fellow researchers - Ashish Garg and Ragini Rahalkar, presented a paper in the 2006 IEEE Information Assurance Workshop guide entitled “Profiling Users in GUI Based Systems for Masquerade Detection” [21]. In this document, the team introduced a new framework that created an individual feature set for a user’s behavior on GUI systems. They collected real user behavior data from live systems and removed limitations to create feature vectors. These vectors contained user information such as mouse angles, speed, and the number of clicks during a user session. For this purpose, they developed an active system logger with C programming and a Microsoft .NET framework. The logger was designed to collect user inquiries on a system in real-time. The team shaped their method of user identification and masquerade detection as a *binary classification problem* and applied a Support Vector Machine (SVM) to categorize these vectors. Their technique provided up to 96% accuracy in user information and detection rates [21].

Kwiat, Upadhyaya, Garg, and Rahalkar discussed related work in the field of masquerade detection such as the application of UNIX; in which user activity is monitored and recorded as the command sequence. The team explains feature extraction details such as data collection, calculations of features, and the use of SVM for learning and classifications. The group presented test results proving their accuracy rate of over 96%. It detailed with 8 features, a 73.85% detection rate was calculated with 68 false positives. However, with 16 features, the detection rate reached 96.15% with only 37 false positives. By increasing the number of features utilized, the team was able to make more accurate distinctions between multiple users. The team concluded that user behavior features based on mouse activity on a GUI system uniquely identified users and thus provided better masquerade detection capabilities [21].

14. Loss Interference in Networking

In 2007, Vidyaraman Sankaranarayanan of the University at Buffalo, Kwiat, and Upadhyaya published a paper on a Quality of Service (QoS) loss interference model. Entitled “QoS-LI: QoS Loss Interference in Disadvantaged Networks” [22], the paper describes how the end nodes, communicating over a disadvantaged network, can employ the model in order to deduce the nature of a QoS loss in a non-intrusive manner. The goal of capturing the nature of QoS loss led them to assume that such loss was logically equivalent to message loss. By developing a game-theoretic concept that captures the nature of message loss they could likewise discern the source of QoS degradation. Their studies included investigations into the integration of the framework into existing platforms. They summarized the problems that involved link selection (as opposed to route selection) in disadvantaged networks as a selected resource problem. The team applied a game-theoretical model to set limits on the rate of game convergence. Applying this convergence rate allowed the loss inference model to distinguish between adversarial network exploitation and benign network loss. The model was designed to help manage the operation of defending military networks by enabling appropriate action to be taken based on the knowledge that the communications were facing a denial-of-service attack. The model provided an appropriate level of verification that the communications were not merely experiencing a period of statistically reasonable message loss and that recovery would likely ensue.

This same period also marked a similar QoS-intensive step for the Assured Communications Research Center. Dr. Kevin Kwiat partnered with graduate student Wenjing Wang and Professor Mainak Chatterjee from the University of Central Florida to research network strategies and user inputs. The three published papers that dealt with a proposed, simple user interface for translation of a user’s inputs into networking parameters that are used by various layers of the protocol stack. One paper was published in the 2007 International Symposium on Wireless Pervasive Computing [23] and the other was published in the 2007 Med Hoc Net Conference [24]. They proposed that end users often find it difficult to express their desired QoS in technical parameters, even when they have the desired QoS in mind. Kwiat, Wang, and Chatterjee conducted a case study where they considered the routing strategies in an ad hoc network environment. They displayed the translation of user preferences into network parameters at the

routing layer. Depending on the user inputs, the most suitable routing strategy was selected to deliver the desired QoS. Simulation results demonstrated how the QoS management affected the network performance via the proposed interface and substantiated their notion of a QoS Dashboard's functionality. Network simulator (ns-2) based results also demonstrated how delay, throughput, and network lifetime were affected when the routing strategy was changed [24].

The three researchers continued their research in ad hoc networks by developing a model of the network as a directed weighted graph, and focusing on the problem of routing in resource rationed ad hoc networks. Their approach began with a hierarchical bidding process, through which nodes in the network hid for virtual currency. The winning bids utilized the granted virtual currency to pay the intermediate nodes for packet forwarding. They studied the trustworthiness of hierarchical routing with rational nodes as well as how network-based security mechanisms could be called upon to identify the malicious nodes. They showed performance gains in data throughput and energy consumption in the ns-2 simulations. The results of this work were likewise published [25].

In other work related to the combined treatment of QoS and security, the ACRC proposed a mathematical integrity function to study how to integrate computer and network components together to build safer systems. Once these systems are composed, they are placed in an adversarial environment. Together with Professor Sibabrata Ray and graduate student Qunwei Zheng from the University of Alabama, Kwiat applied their integrity function to devise strategies for resilient server replication and placement in a hostile environment [26]. The team devised polynomial-time algorithms to compute the integrity function under benign fault models. They also computed these algorithms under Byzantine fault models, which provided corresponding heuristics. Their goal was to provide algorithms that provided means to compare replication and placement strategies against intelligent adversaries. They also described how these algorithms could aid in computing good replication and placement in adversarial environments. Optimal resource allocations were feasible and would satisfy the system's performance criteria while staying within the system's resource limitations.

In 2008, Vidyaraman Sankaranarayanan, Kwiat, and Upadhyaya renewed their pursuit of inferring the nature of QoS loss in networks when they published “QoS-LI: QoS Loss Interference in Disadvantaged Networks – Part II” [27]. This publication focused on the translation of the theoretical framework to a satellite-based disadvantaged network. When experiencing QoS loss caused by statistical variations, it is possible for the communicating end-points to engage in logical-link switching mode in order to return to the former assured service quality level. However, their game theoretic model revealed the expected: in the presence of a jammer, the return to a previously assured QoS level is expected to take more time and the end-to-end delay is not expected to rapidly improve. Future work in this area will investigate the use of metrics other than end-to-end delay to infer QoS loss, and when multiple streams of data are sent to the across satellite links to a base station.

Their results strengthened their earlier work by integrating the model with known operational parameters used in satellite communications. Using OPNET simulations, they were able to demonstrate the feasibility of integrating their game theoretic formulation in real disadvantaged networks. The operational parameters used in this work were open-source information. For example, the OPNET simulations used operational parameters that were already built into the simulation package.

15. ACRC Research in Wireless Networks

At the Fourth International Conference on Mathematical Methods, Models and Architectures for Computer Network Security, Upadhyahya, Kwiat, et al, presented their paper entitled “Environment- Aware Trusted Data Delivery in Multipath Wireless Protocols” [28]. They proposed a trust-based route selection framework for Multi-Hop Wireless Networks (MWNs). The team presented a modification to the Ad-hoc On-demand Multipath Distance Vector (AOMDV) protocol. They modeled the AOMDV protocol to include metrics that can be adjusted to varying network conditions when calculating a packet's paths.

Their model maintains two trust values: *Route Trust* values that measure the reliability of packets to reach destination on a specified path, computed by each node for all routes in routing table; and *Node Trust* values that measures the confidence of one-hop neighbors then forwards their measures towards the destination nodes. The Node Trust value is initialized at 1 for destinations and .5 for all other participating nodes. Node trust values correspond to the performance of nodes downstream (in the direction toward the destination). Route Trust values are made up for the Effective Link Capacity (ELC) that indicates traffic on a link for a particular destination and the Effective Route Capacity (ERC) value which recursively factors the ELC into a function that computes the effective capacity of a route from an intermediate node to the destination route.

Traditionally, the source node broadcasts a Route Request (RREQ) packet that is forwarded by neighbor nodes until a route to the destination is discovered. Once a node receives the request, the nodes respond with a unicast Route Reply (RREP) back to the source node with updated information on links and nodes. The team modeled these RREQ and RREP packets to include ELC and ERC values. The ELC and ERC values are used to build a data structure called the Neighbors' Trust Table that contains neighbor node ID's and corresponding trust values. The trust table is updated every time a node receives a RREP packet. In the event a node wants to reassess its route trust, it sends a Query (QRY) packet to a destination. The destination replies with a Query Acknowledgement (QRY-ACK) packet. The source and intermediate nodes update their trust tables once it receives the QRY-ACK packet.

Packet delivery ratio and trust convergence latency were evaluated through GLOMOSIM 2.02 simulations. The team evaluated a variety of scenarios: general congestion, simultaneous congestions on single and multiples routes, etc. Their routing scheme's effectiveness establishes a cyber defense, self-learning network that adapts to environment conditions.

16. Overlaid Wireless Access Networks in Market-Based Environment

With the increasing demand for wireless services, there is a pressing need for the wireless service providers to use different access technologies to harness different levels of coverage, bandwidth, and reliability, so as to provide consistent quality of service at competitive prices. In 2007, Shamik Sengupta of Stevens Institute of Technology, AFRL visiting professor Mainak Chatterjee, and Kwiat proposed the use overlaid networks where a single service provider tries to maximize its revenue by combining heterogeneous access technologies to offer a desired QoS [29].

Their market-based model captures the interaction between providers and users in a competitive environment where users had the option to choose who they perceive as the best service provider. Once a user chooses the best service provider based on the perceived utility, the service provider decides on the selection of the best network to cater to the user's service while maintaining the QoS of the existing users. They conducted a simulation-based case study where a wireless service provider used two overlapping networks: a 3G network that provided a larger coverage area but with low bandwidth and an IEEE 802.11 Wi-Fi network that had a higher bandwidth but a smaller coverage area. The simulations showed the performance of the proposed schemes by measuring the per-user utility and perceived bandwidth for increasing cell radius with both ideally and randomly deployed Wi-Fi access points. The results demonstrated the proposed network selection mechanism achieved better net utility for both the providers and users.

17. Maximizing Spectrum Utilization in IEEE 802.22

The IEEE 802.22 working group is currently developing a Wireless Regional Area Network (WRAN) technology that aims at efficient use and reuse of white space in the TV allocated frequency spectrum. In May 2008, Shamik Sengupta, Mainak Chatterjee and Kevin Kwiat presented a paper entitled “Interference Aware Spectrum Allocation in IEEE 802.22 Wireless Mesh Networks” [30], at the Wireless and Optical Communication (WOC) conference. Their paper examines the current IEEE 802.22 system architecture and its limitations in creating wireless back-haul mesh networks. The goal of their research was to develop dynamic spectrum access among devices in the mesh network. The original ACRC goal of adapting fault tolerance for cyber defense is related to spectrum access: making mesh architectures attack-tolerant calls for investigating new approaches to assuring the communications between the mesh’s nodes.

Segupta, Chatterjee and Kwiat proposed a coordinated distributed scheme for efficient spectrum allocation that increases utilization and reduces inference among nodes that compete for access to the spectrum. In their proposed scheme, they team employed a point-to-point (mesh) wireless regional access network (WRAN) topology and cast the access problem in graph theoretic terms. Once the graph model for the problem was created, they employed their Maximum Utility Graphing Coloring (MUGC) algorithm. MUGC dynamically allocates the spectrum for a mesh network, which in turn enables higher spectrum utilization with fewer collisions. Simulation-based experiments demonstrated that the proposed MUGC algorithm outperforms standard spectrum allocation.

18. Education for Potential Air Force Officers

In 2003, the Air Force Research Laboratory initiated the United States Air Force Advanced Course in Engineering (ACE) Cyber Security Boot Camp. This 11-weeks summer program has provided cyber security education for Reserve Officer Training Corps (ROTC) cadets and university students from all over the nation. This highly selective, highly challenging course provides students with the knowledge and the education in critical areas of cyberspace such as malware, digital forensics, and cyber warfare [31]. Each summer, selected ACE students join the ACRC to perform studies in the history of fault-tolerant computing, attack-resistant replication, and distributed voting algorithms. The work performed by the students has been published in the newsletter of the DoD Information Assurance Technology Analysis Center (IATAC) newsletter – *The IANewsletter* [32]. Their research and educational development with the ACRC emphasizes the ACE goal to extend beyond mere academic expectations by motivating these interns to potentially become future personnel in cyber security.

19. Summary

Initially, the ACRC pursued computer systems that are both secure and fault-tolerant. From this pursuit emerged several sub-areas of research: intrusion detection, applications in mobile computing, simulation-based validations, coordination models, secure knowledge management, internet worms, QoS loss inference, SCADA systems, GUI-based systems, and wireless systems. The communications requirements that underlie replication drove these pursuits. Now the use of replication, coupled with the ACRC's voting protocols, can become the basis for a "fight through" capability in cyber defense. Most recently, the ACRC's research in voting protocols was published in a journal [33]. We envision our proposed follow-on work of the ACRC will include a more complete explication, as other cyber defense mechanisms take shape, of "fight through."

20. References

1. Benjamin Hardekopf, Kevin Kwiat, "Performance Analysis of an Enhanced-Security Distributed Voting Algorithm," *2000 Symposium on Performance Evaluation of Computer and Telecommunication Systems*, Simulation Councils, Inc., 2000, Vancouver, British Columbia, pp. 342-351.
2. Benjamin Hardekopf, Kevin Kwiat, *Distributed Voting for Security and Fault-Tolerance*, May 2001 In-House Report: Air Force Research Laboratory, Information Directorate, Rome, NY: AFRL-IF-RS-TR-2001-53.
3. Kaliappa Ravindran, Kevin Kwiat, Ali Sabbir, "Adapting Distributed Voting Algorithms for Secure Real-Time Embedded Systems," *The 24th IEEE International Conference on Distributed Computing Systems Workshops*, 23-24 Mar 2004, Hachioji, Tokyo, Japan, pp.347-353.
4. Kaliappa Ravindran, Kevin Kwiat, Ali Sabbir, "Incorporating Timeliness in Atomic Write of Shared Data for Distributed Embedded Systems," *The 24th IEEE International Conference on Distributed Computing Systems Workshops*, 23-24 Mar 2004, Hachioji, Tokyo, Japan, pp. 884-889.
5. Jiang Wu, K. Ravindran, K. Kwiat, A. Sabbir, "Adaptive Voting Algorithms for Reliable Dissemination of Data in Sensor Networks," *Proceedings of the International Conference on Availability, Security and Reliability*, 4-7 March 2008, Barcelona Spain
6. Kevin Kwiat, Shambhu Upadhyaya, "A Comprehensive Simulation Platform for Intrusion Detection in Distributed Systems," *The Proceedings of the 2000 Summer Computer Simulation Conference*, Simulation Councils, Inc., 2000, Vancouver, British Columbia, pp. 586-591.
7. Ramkumar Chinchani, Shambhu Upadhyaya, Kevin Kwiat, "Towards the Scalable Implementation of an Anomaly Detection System," *Proceedings of the Military Communications Conference*, (MILCOM) 2002, Anaheim, CA, USA 2002.
8. Charles Choi, "Computer Security: Keyboard Cops," *Scientific American*, December 2002.
9. Ramkumar Chinchani, Shambhu Upadhyaya, Kevin Kwiat, "A Tamper-Resistant Framework for Unambiguous Detection of Attacks in User Space Using Process Monitors," *IEEE International Information Assurance Workshop*, 2003, Darmstadt, Germany, pp. 25-33.

10. Ramkumar Chinchani, "A Job-Centric Approach to User-Level Intrusion Detection," Doctoral Dissertation, Graduate School of the State University of New York at Buffalo, 2005, Buffalo, NY.
11. Shambhu Upadhyaya, Ramkumar Chinchani, Kiran Mantha and Kevin Kwiat, "Encapsulation of User's Intent: A New Proactive Intrusion Assessment Paradigm," (book chapter) in *Managing Cyber Threats: Issues, Approaches and Challenges*, Kluwer Academic Publishers, 2005.
12. Zesheng Chen, Lixin Gao, Kevin A. Kwiat, "Modeling the Spread of Active Worms," *Proceedings of INFOCOM*, 2003, San Francisco.
13. Jiang Wu, Sarma Vangala, Lixin Gao, and Kevin A. Kwiat, "An Effective Architecture and Algorithm for Detecting Worms with Various Scan," *Proceedings of the Internet Society Network and Distributed System Security Symposium (NDSS)*, 2004, San Diego.
14. Jianhong Xia, Sarma Vangala, Jiang Wu, Lixin Gao, and Kevin A. Kwiat, "Effective Worm Detection for Various Scan Techniques," *Journal of Computer Security* Vol. 14, No. 4, 2006, pp. 359-387.
15. Kaliappa Ravindran, Kevin Kwiat, G. Ding, "Simulation-based Validations of Protocols for Concurrent Systems," *2004 IEEE Global Telecommunications Conference Workshops*, 29 Nov – 3 Dec 2004, Dallas, Texas, pp. 331-340.
16. Kevin Kwiat, Shangping Ren, "A Coordination Model for Improving Software System Attack-tolerance and Survivability in Open Hostile Environments," *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing*, 2006, Taichung, Taiwan, pp. 394-402.
17. Shangping Ren, Yue Yu, Kevin Kwiat, Jeffrey Tsai, "A Coordination Model for Improving Software System Attack-Tolerance and Survivability in Open Hostile Environments," *International Journal of Distributed Sensor Networks*, Vol. 3, Issue 2, 2007, Philadelphia, PA, pp. 175-200.
18. Yu Yue, Shangping Ren, Kevin Kwiat, "Take Intelligent Risk and Optimize Decision Based on Time, Available Resources and Tolerance Limits," *Proceedings of the 13th IEEE Real Time and Embedded Technology and Applications Symposium*, 3-6 Apr 2007, Bellevue, Washington, pp. 315-325.
19. Kun Xiao, Nianan Chen, Shangping Ren, Limin Shen, Xianhe Sun, Kevin Kwiat, Michael Macalik, "A Workflow-based Non-intrusive Approach for Enhancing the Survivability of Critical Infrastructures in Cyber Environment," *Proceedings of the 2007 International Conference on Software Engineering (ICSE)*, 20-26 May 2007, Minneapolis, MN.

20. Kun Xiao, Shangping Ren, Kevin Kwiat, "Retrofitting Cyber Physical Systems for Survivability through External Coordination," *41st Hawaii International Conference on System Sciences*, 7-10 January 2008, Waikoloa, Big Island, Hawaii.
21. Ashish Garg, Ragini Rahalkar, Shambhu Upadhyaya, Kevin Kwiat, "Profiling Users in GUI Based Systems for Masquerade Detection," *2006 IEEE Information Assurance Workshop*, 21-23 June 2006, West Point, New York, pp. 48-54.
22. Vidyaraman Sankaranarayanan, Kevin Kwiat, Shambhu Upadhyaya, "QoS-LI: QoS Loss Inference in Disadvantaged Networks," *IEEE Computer Society 21st International Conference on Advanced Information Networking and Applications Workshops/Symposia*, 2007, Los Alamitos, CA, pp. 524-529.
23. Wenjing Wang, Mainak Chatterjee, Kevin Kwiat, "QoS Dashboard: Translation of User Inputs to Networking Strategies," *2007 International Symposium on Wireless Pervasive Computing*, 5-7 Feb 2007, San Juan, Puerto Rico, pp. 20-24.
24. Wenjing Wang, Mainak Chatterjee, Kevin Kwiat, "Human-Network Interface for QoS Management in Ad hoc Networks," *2007 Med Hoc Net 6th Annual Mediterranean Ad Hoc Networking Workshop*, 13-15 June 2007, Corfu, Greece, pp. 17-24.
25. Wenjing Wang, Mainak Chatterjee, Kevin Kwiat, "An Economic Approach to Hierarchical Routing in Resource Rationed Ad Hoc Networks," *8th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, 18-21 June 2007, Helsinki.
26. Sibabrata Ray, Qunwei Zheng, Kevin Kwiat, "Energy-efficient distributed Authorization Server (DAS) Placement for Sensor Networks," *The International Journal of High Performance Computing and Networking*, Vol. 1, No.4, Inderscience, 2005.
27. Vidyaraman Sankaranarayanan, Kevin Kwiat, Shambhu Upadhyaya, "QoS-LI: QoS Loss Interference in Disadvantaged Networks – Part II," *Proceedings of the 11th Communications and Networking Simulation Symposium (CNS)*, ACM/SIGIM, 14-17 Apr 2008, Ottawa, Canada.
28. Mohit Virenda, Arunn Krishnamurthy, Krishnan Narayanan, Shambhu Upadhyaya, and Kevin Kwiat, "Environment-Aware Trusted Data Delivery in Multipath Wireless Protocols," *Proceedings of the 4th International Conference on Mathematical Methods, Models, and Architectures for Computer Network Security (MMM-ACNS)*, 13-15 Sept 2007, St. Petersburg, Russia.
29. Shamik Sengupta, Mainak Chatterjee, and Kevin Kwiat, "Pricing-Based Service and Network Selection in Overlaid Access Networks", *Proceedings of the 6th International Conference on Information, Communications, and Signal Processing*, IEEE 10-13 Dec 2007, Singapore.

30. Shamik Sengupta, Mainak Chatterjee, and Kevin Kwiat, "Interference Aware Spectrum Allocation in IEEE 802.22 Wireless Mesh Networks," *Proceedings of the 8th International Conference on Wireless and Optical Communications (WOC)*, IASTED, 26-28 May 2008, Quebec City, Quebec, Canada.
31. "Advanced Course in Engineering Cyber Security Boot Camp,"
<https://www.acecybersecurity.com/>
32. Kevin Kwiat, Shambhu Upadhyaya, and Amber Helton, "A Decade of Air Force and Academic Collaboration Toward Assuring Information," *IAnewsletter: The Newsletter for Information Assurance Technology Professionals*, Vol. 10, No. 3, DoD Information Assurance Technology Analysis Center (IATAC), Fall 2007.
33. Kaliappa Ravindran, Kevin Kwiat, and Patrick Hurley, "Adaptive Voting Algorithms for the Reliable Dissemination of Data in Fault-Prone Distributed Environments," *International Journal of Business Intelligence and Data Mining*, Vol. 3, No. 3, 2008.